

CAS Production Line by Vijay Sekhri



Ideal Production Line

- Excellent budget
- Good availability of hardware
- Manpower and commitment
- Maximum output through parallelism
- Move Data from one machine onto another

Data Collection

Collects data from sky
And stores in FITS files

- Raw data is received at Fermilab from APO, processed, calibrated, and stored in binary FITS files.
- These binary FITS files are stored on the Data Archive Server (DAS) for collaboration and public access.
- All file handling and storage to this point occurs on Linux boxes.

Data Collection


Collects data from sky
And stores in FITS files



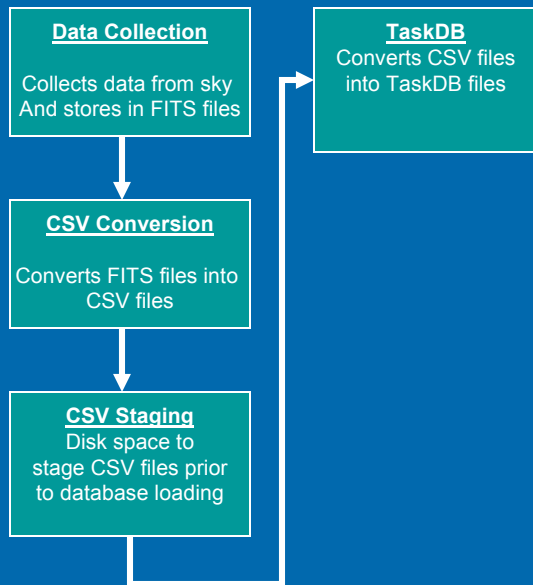
CSV Conversion

Converts FITS files into
CSV files

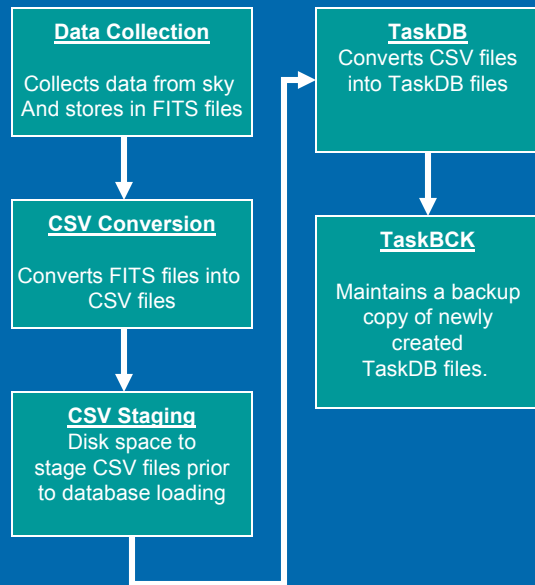
- Convert the FITS files into CSV files on Linux Box
- CSV files are generated for spectroscopic data (spectro CSVs) and two versions of imaging data (Target and Best).



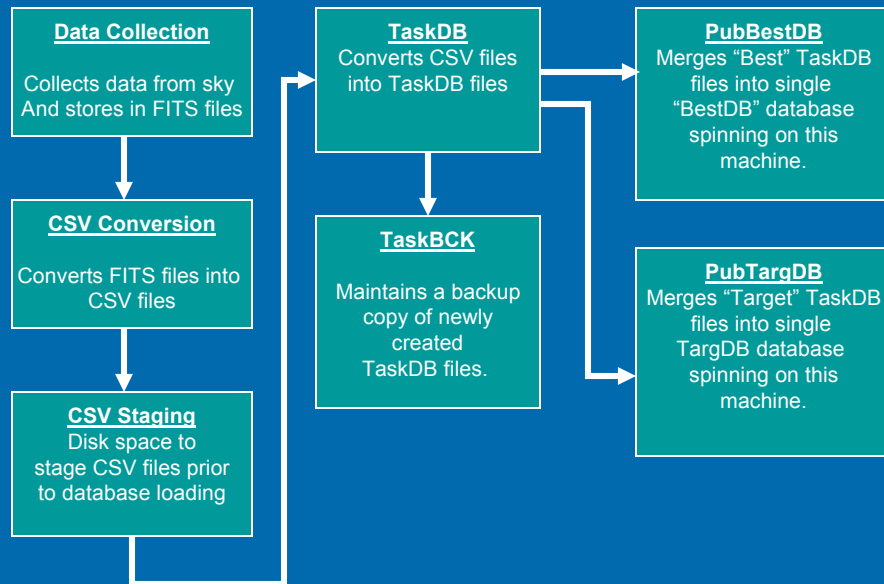
- Copy the CSV files onto a Windows Machine
- Not enough disk space left on Linux machines. Make a backup of older CSV files that have been successfully "published" and then delete the copy spinning on disk to free up space for subsequent CSVs.
- Availability of this machine should be approx 95%



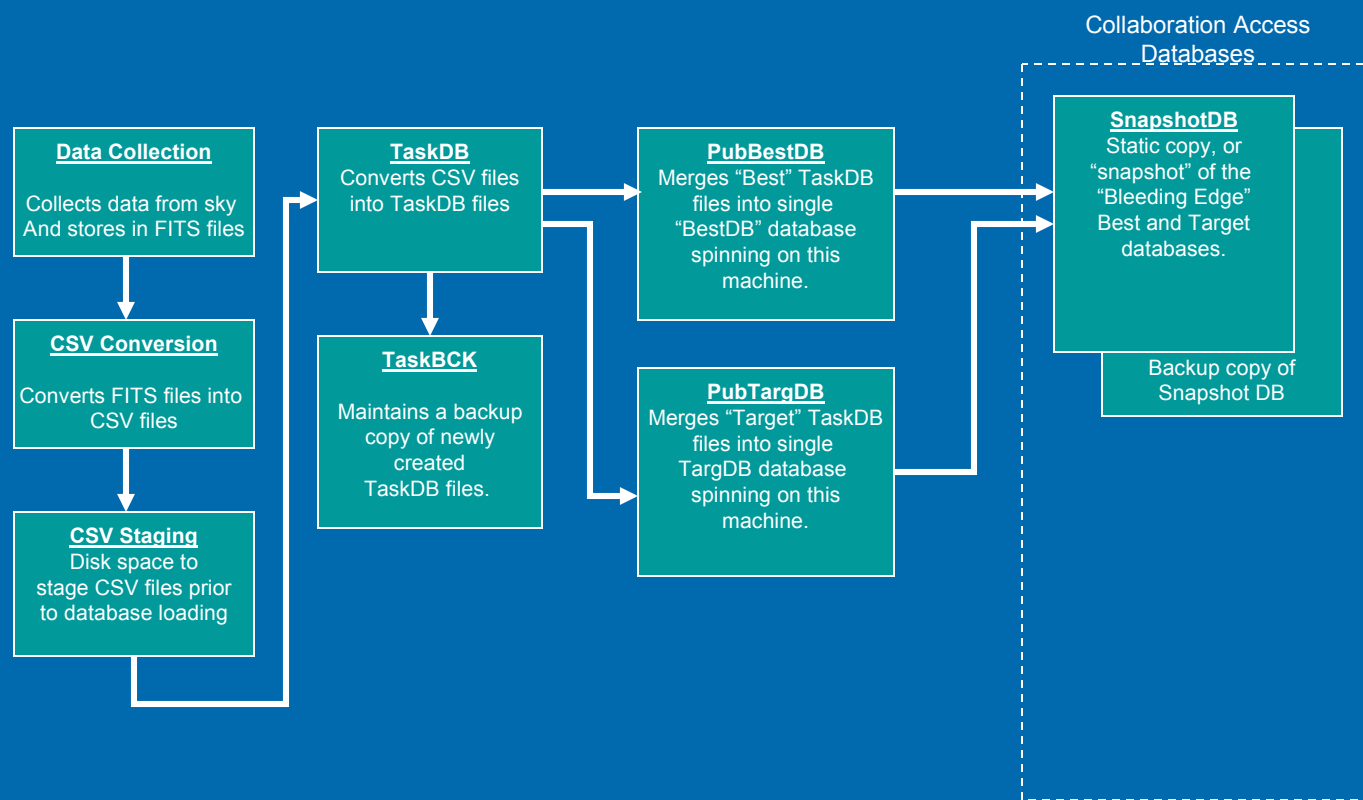
- Data access to the collaboration and public is provided through two interfaces: the DAS and the CAS. The CAS contains the full object catalog and enables users to submit complex queries using output parameters from the data reduction pipelines and calibration processes.
- sqlLoader loads CSV into MSSQL Server 2000
- Each CSV correspond to one Task on MSSQL Server
- Each Task generates a corresponding database called a TaskDB
- Target, Best and Spectro TaskDB are produced.
- Performs validation and accuracy check of CSV file
- If errors are detected, we troubleshoot upstream, fix the problem, and re-generate the CSV file.



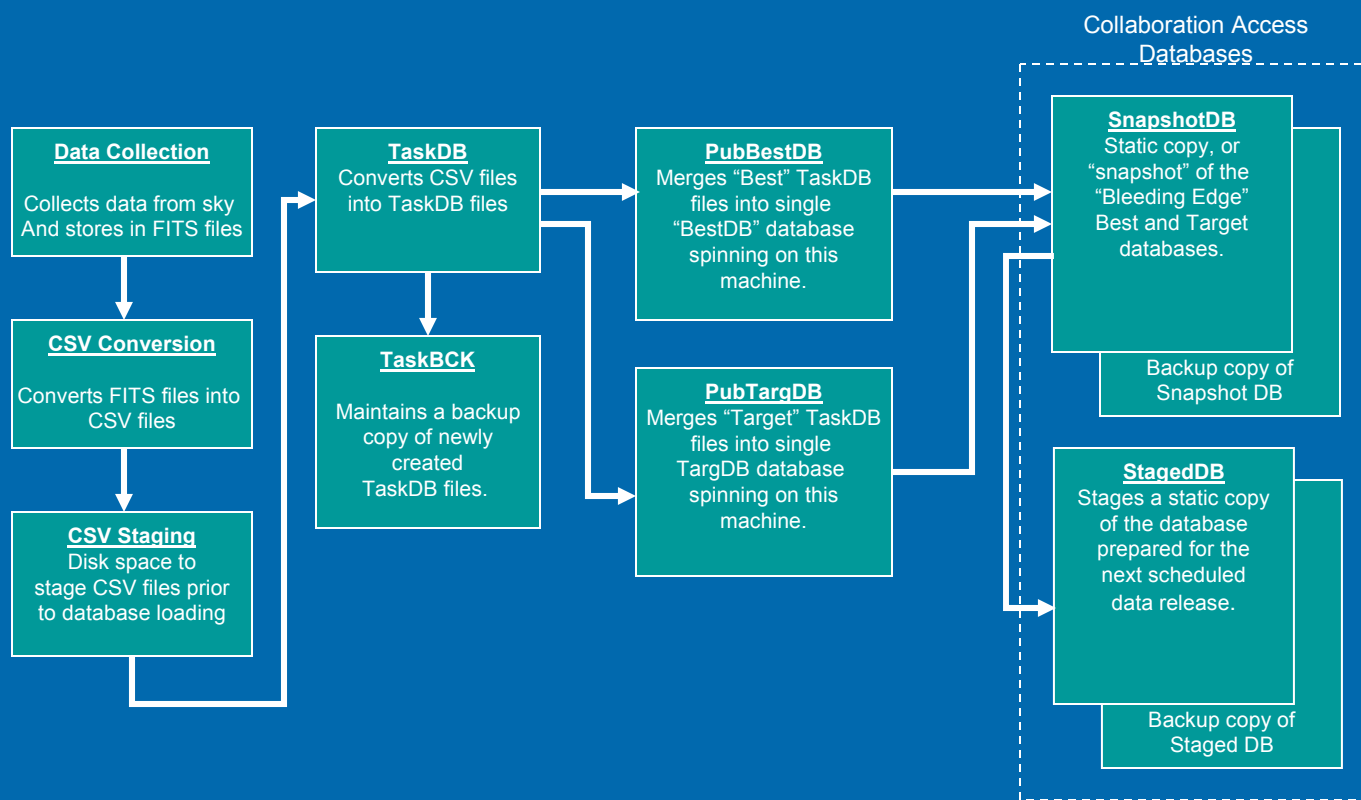
- Parallel step with TaskDB
- Back up each TaskDB onto a different Windows machine, and perhaps eventually to tape.
- we need to save all TaskDBs to allow for prompt re-creation of a finished database should we have to re-process some fraction of the data.
- At this stage, we could start generating CSV files using data to be loaded for the next data release.



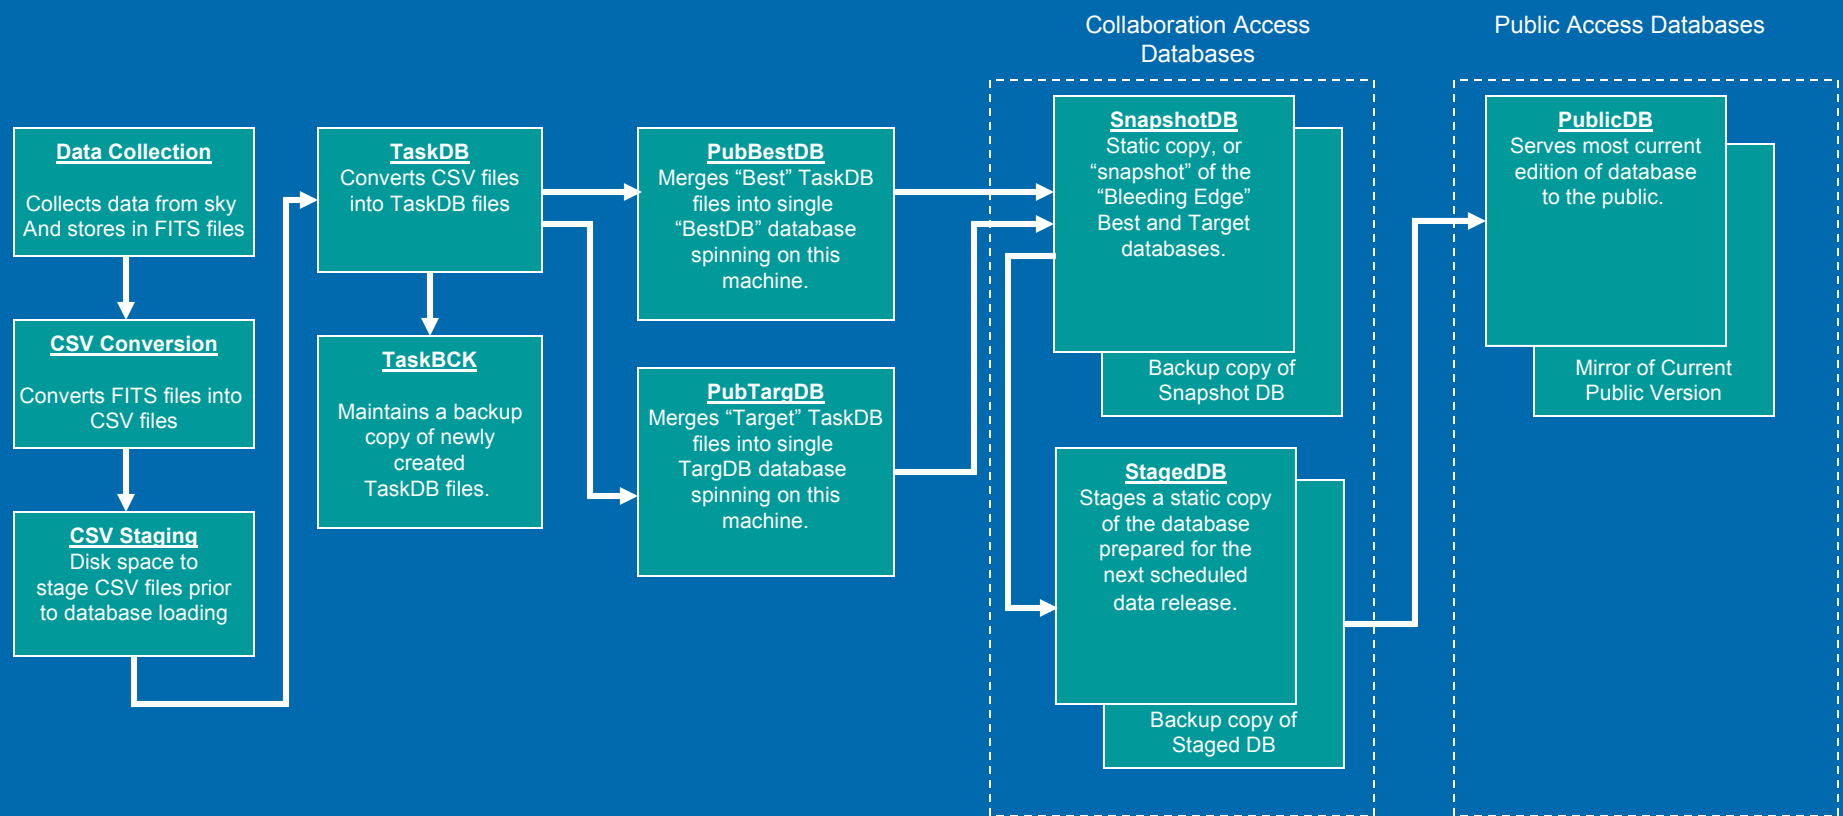
- Combines all the Best and Target TaskDBs into one big PUBBEST and PUBTARG DB respectively.
- Primary keys, secondary keys and neighbors tables are generated for faster search
- This machine is CPU intensive machine and availability should be 95%
- Might need 2 Machines here if the DBs gets too big. Estimated sizes are DR3: 3.9 TB, DR4: 5.1TB Final DR5: 6.1 TB
- Current max size of a machine after format becomes 4.8TB
- At this stage we can start CVS staging if the CVS Conversion is done for next Data Release. Can delete old data or buy new machine.



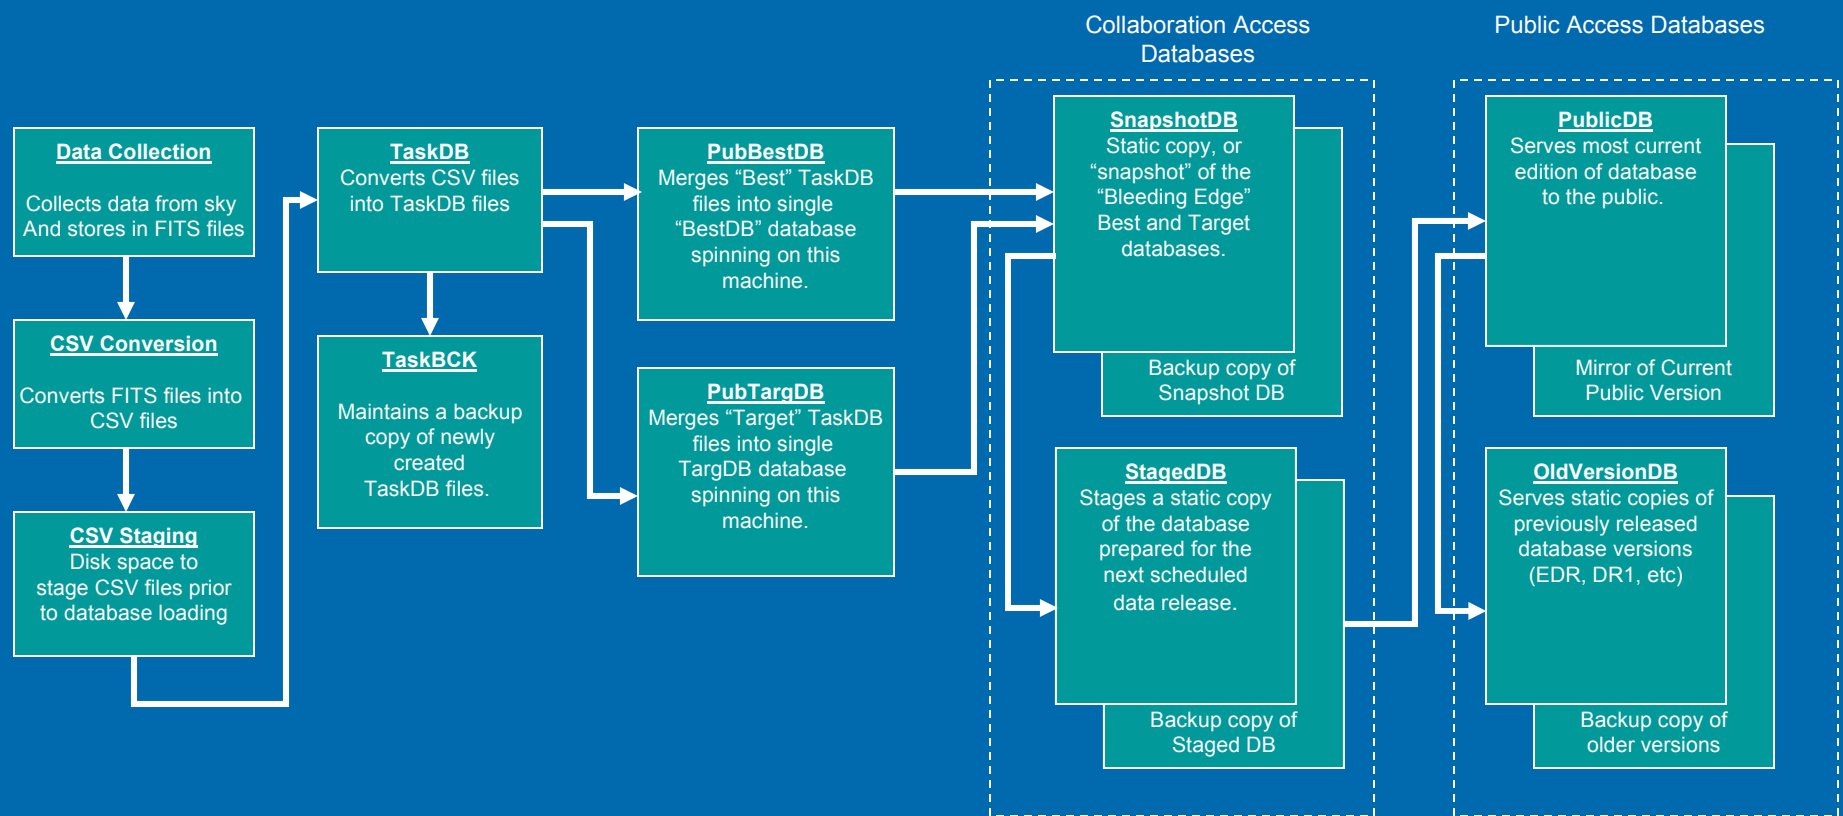
- Detach, copy and attach the PUB DBs on another Windows Machine
- Give Collaboration access to this new DB
- Test the DB for any discrepancies. Troubleshoot and fix any DB-related or upstream problems.
- The "snapshot" copy of the DB will not have the complete set of data for the next release and is subject to change as more data is loaded or as problems are found and corrected. Nonetheless, it does provide testers and the collaboration with access to more recently loaded data.
- Should have almost the same space as PubDB machines with availability of 95%



- Gives Collaboration access to the next data release at least three months before the public release. Provides access for research and final testing and evaluation.
- Static copy of the DB that does not change unless serious problems are found.
- All problems or shortcomings that are not corrected are documented on a "caveats" page posted on the release website.
- Ideally Staged DB should be a different machine altogether.
- Backup copy must exist; otherwise, there is only one copy of the DB for the next public release. Without a backup, we would have to regenerate the entire DB if it somehow became corrupted.
- Machine's availability should be 95% but does not need to be cutting edge high performance server.
- At this stage we can start TaskDB creation from the staged CSV files for next Data Release and the CSV Conversion for next to next Data Release. Parallelism can happen only if we have all the machine at our disposal.
- Every release contains new data acquired through a specified point in time (typically June of the preceding year), so all releases beyond DR2 will be incremental releases.



- Gives the public access to the vetted copy of the imaging and spectroscopic data associated with the new release.
- Static data that does not change
- This machine should be a high performance server with 24/7 uptime and 99% availability
- Backup copy should be available for maintaining the uptime
- Mirroring is done with primary cache here and secondary at JHU. If one server goes down public access automatically trips to the secondary cache
- Another machine is physically shipped to JHU which is returned after copying the Public data for mirroring



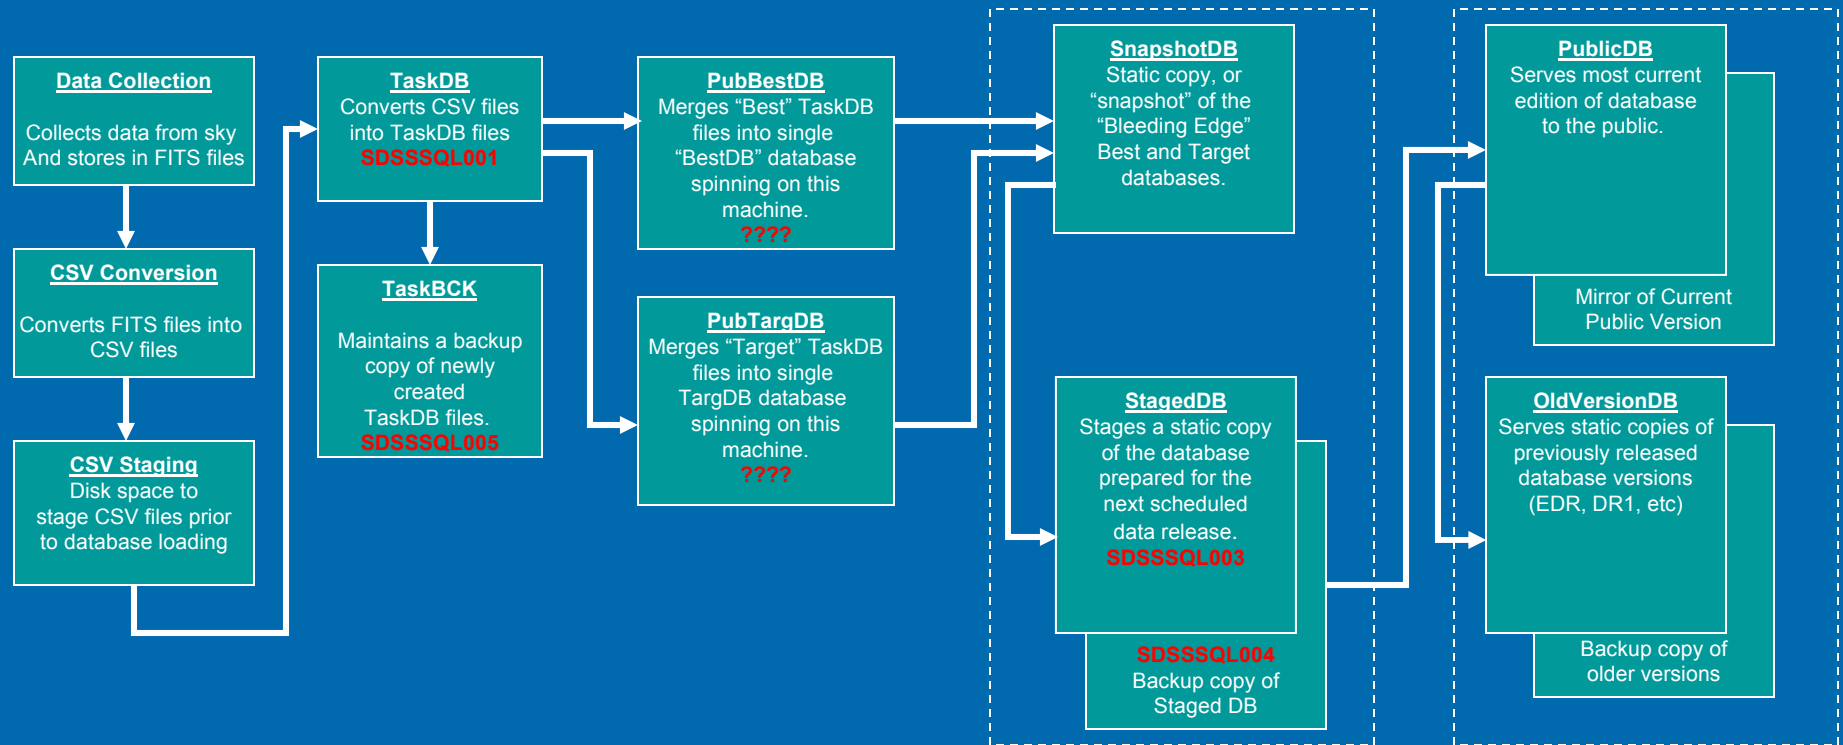
- Every release contains the current data and the data from previous releases
- Older data is pushed down to another machine but still hosted
- At least two previous copies should be maintained for public research
- This machine should be a high performance server with 24/7 uptime and 99% availability. Although availability for prior release can be 95%. This reduces the support load we place on the Computing Division
- Backup copy should be available for maintaining the uptime.

Storage Requirement and Availability

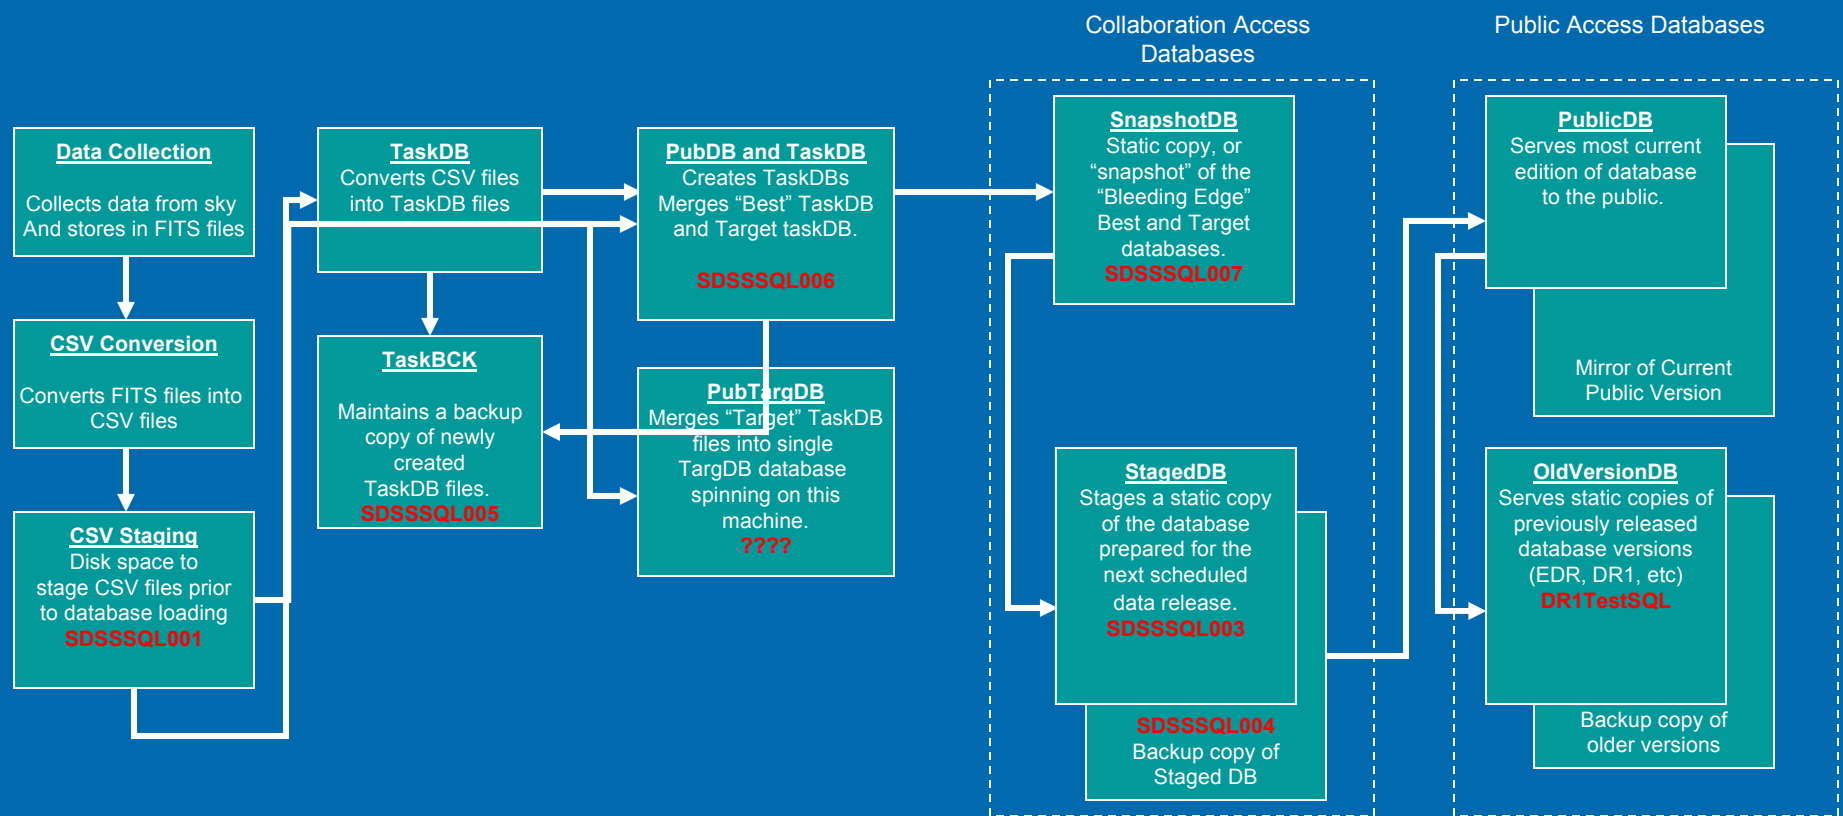
	Network Bandwidth	Usable Disk Space	Availability	Fault Notification/ Acknowledgement Time	Fault Response/ Recovery Time
CSV Staging	1 Gbps	1.5 TB	95%		
TaskDB	1 Gbps	2.0 TB	95%		
TaskBCK	1 Gbps	2.0 TB	99%		
PubBestDB PubTargDB	1 Gbps	DR3: 3.9 TB DR4: 5.1TB Final: 6.1 TB	95%	15 min. / 16 hrs	16 hrs / 48 hrs (no off-hour support)
SnapshotDB	1 Gbps	DR3: 3.0 TB DR4: 4.0 TB Final: 4.8 TB	95%		
StagedDB	1 Gbps	DR3: 3.0 TB DR4: 4.0 TB Final: 4.8 TB	95%	15 min. / 16 hrs	16 hrs / 48 hrs
PublicDB	1 Gbps	DR3: 3.0 TB DR4: 4.0 TB Final: 4.8 TB	99%	Immediately upon switchover to mirror / 30 min.	8 hrs / 24 hrs
OldVersionDB	1 Gbps	DR3: 3.0 TB DR4: 4.8 TB Final: 7.0 TB	95%	15 min. / 16 hrs	16 hrs / 48 hrs (no off-hour support)
Distribution Box	1 Gbps	DR3: 3.0 TB DR4: 4.0 TB Final: 4.8 TB	95%		

Real Production Line

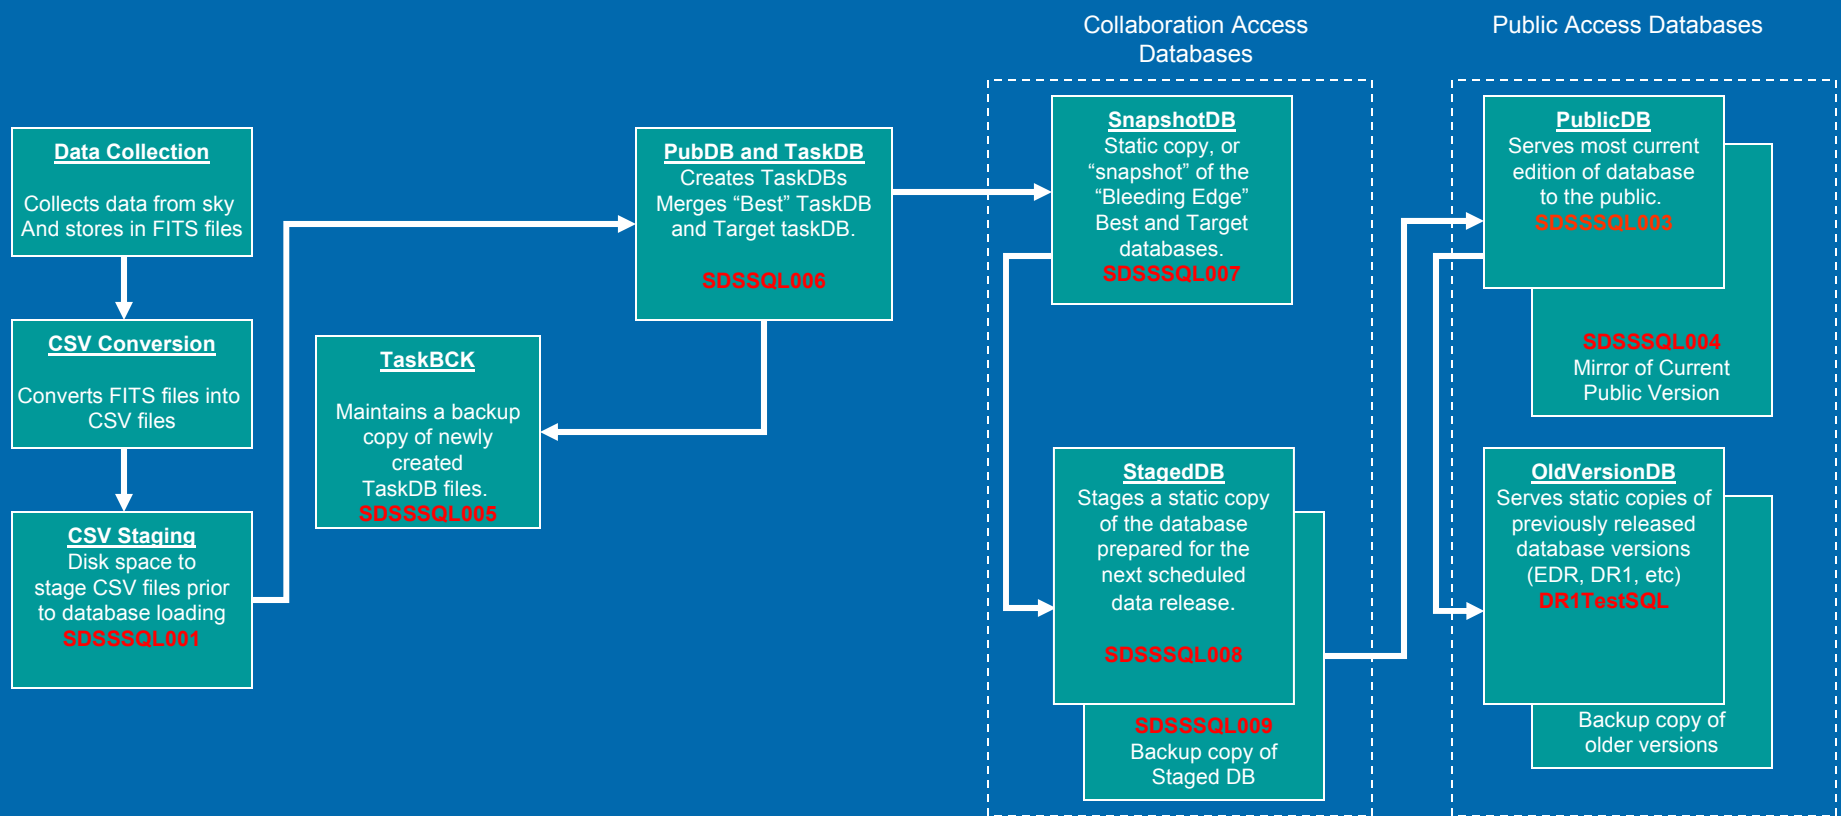
- Limited budget
- Limited Hardware
- Some parallelism through swapping of hardware
- Move machines instead of data for different roles



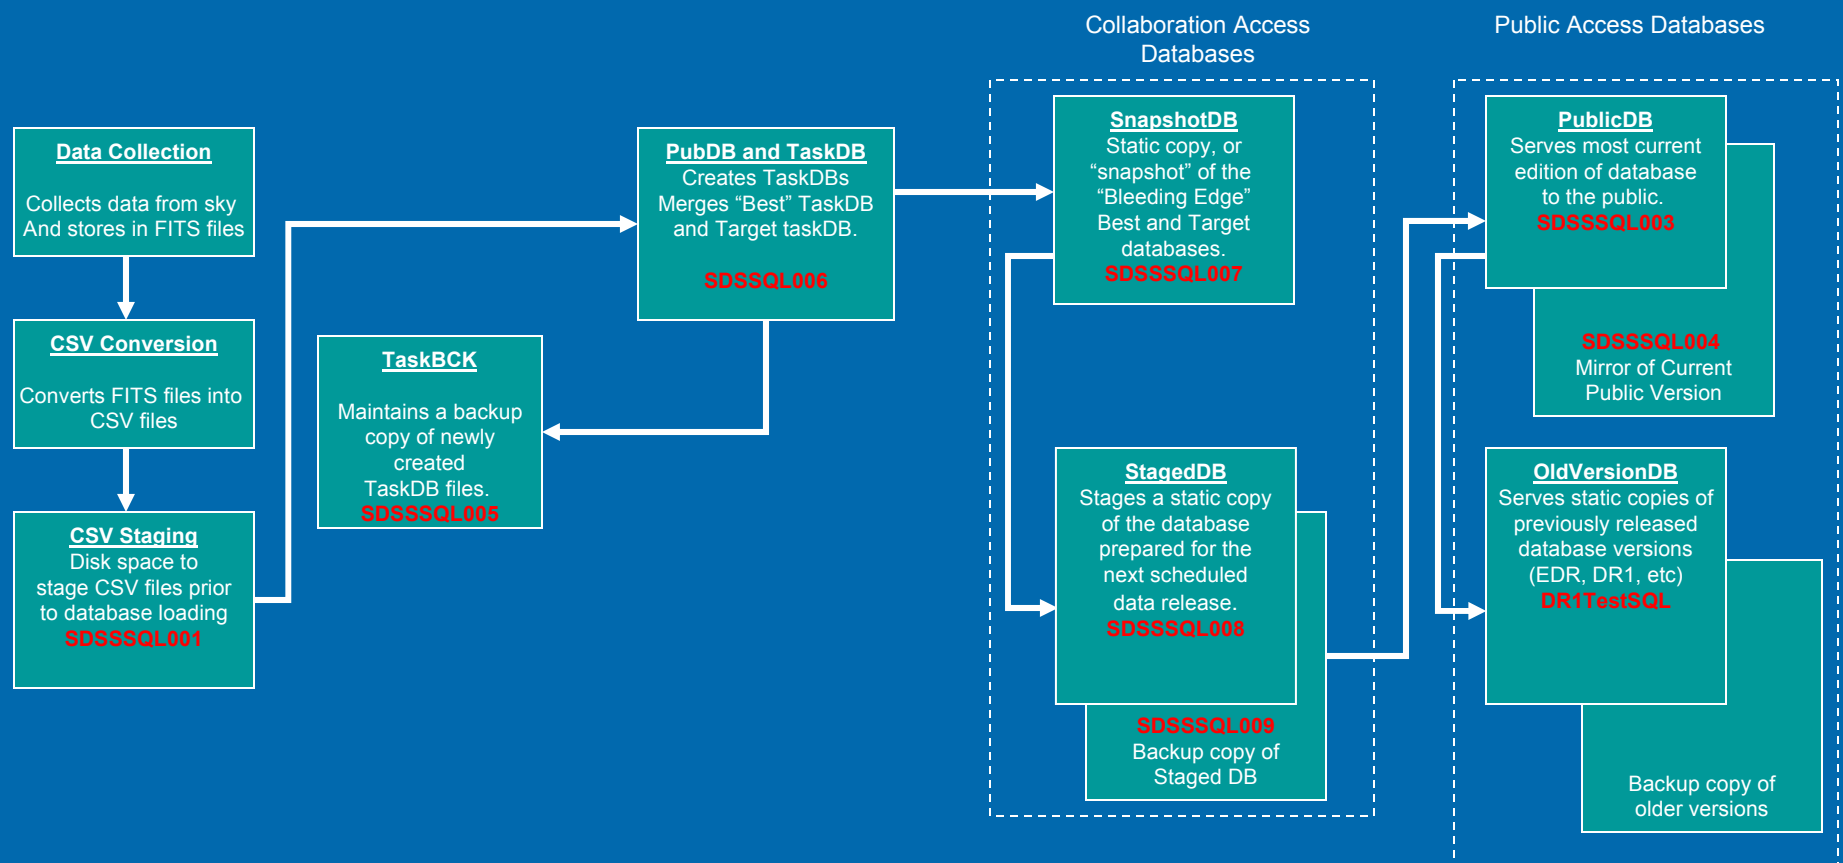
- SDSSSQL003 and SDSSSQL004 are being used for DR2 StagedDB and backup.
- We don't have any backup SnapshotDB.
- 001 and 005 provide storage space for staging CSVs and TaskDBs for the DR3 load.
- 002 is used for web interface to run this production line.



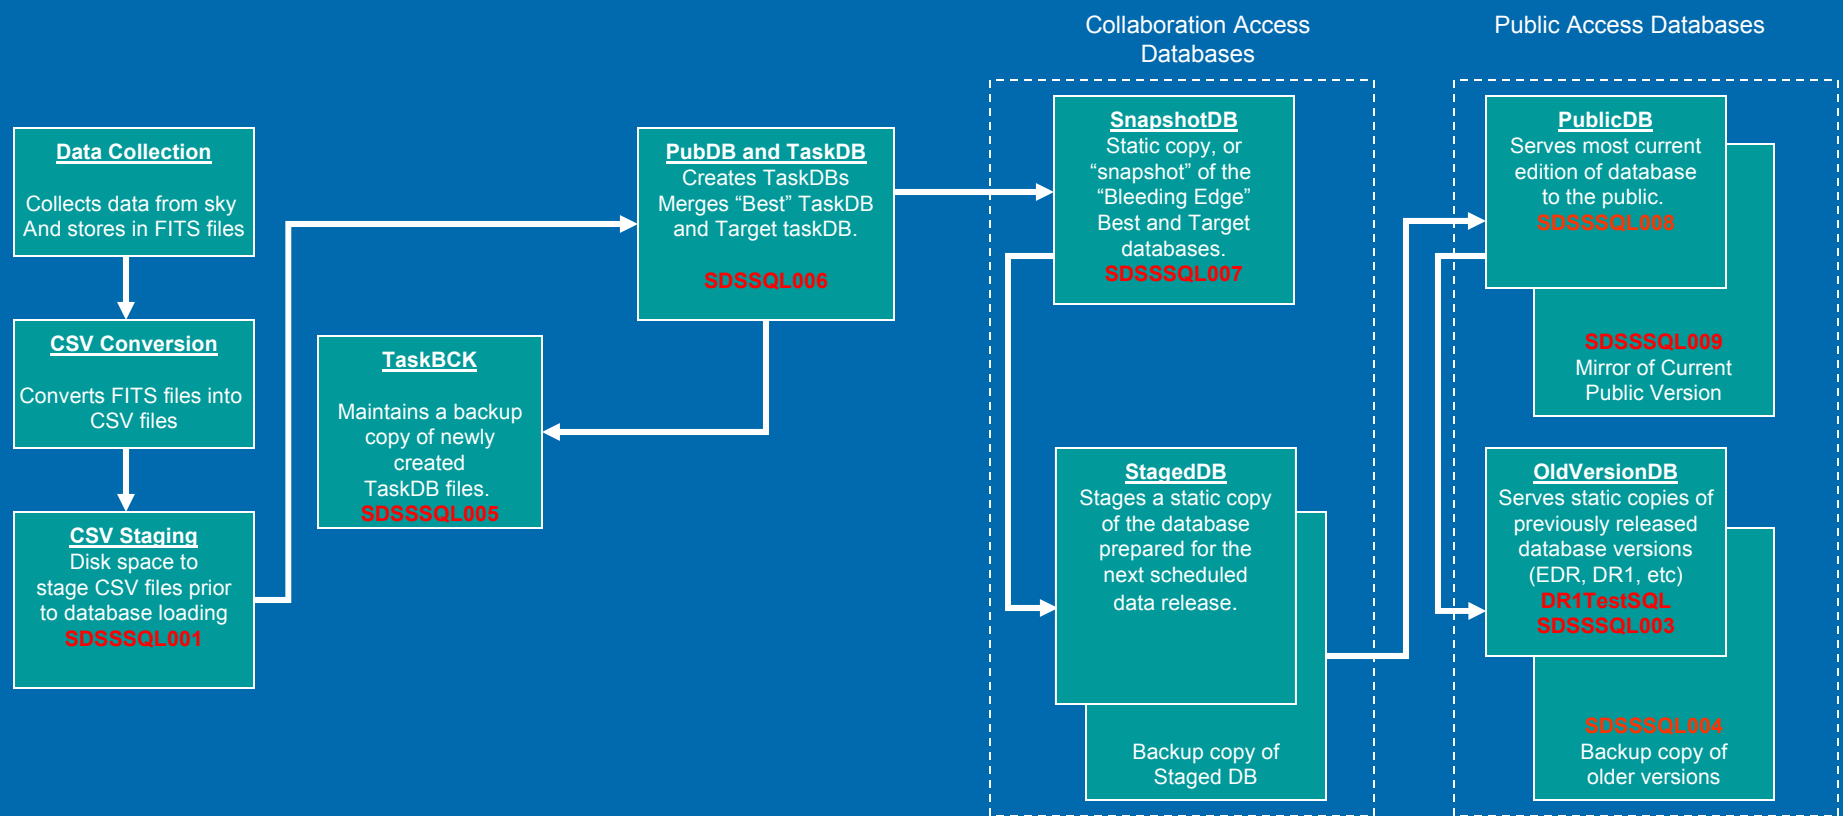
- TaskDB, PubBestDB and PubTargDB becomes a single machine.
- In February 2004, we purchased two machines (006 and 007) to serve as PubDB and SnapshotDB for the DR3 load.
- On March 15, 003 and 004 move down to become the PublicDB and backup for the DR2 public release.



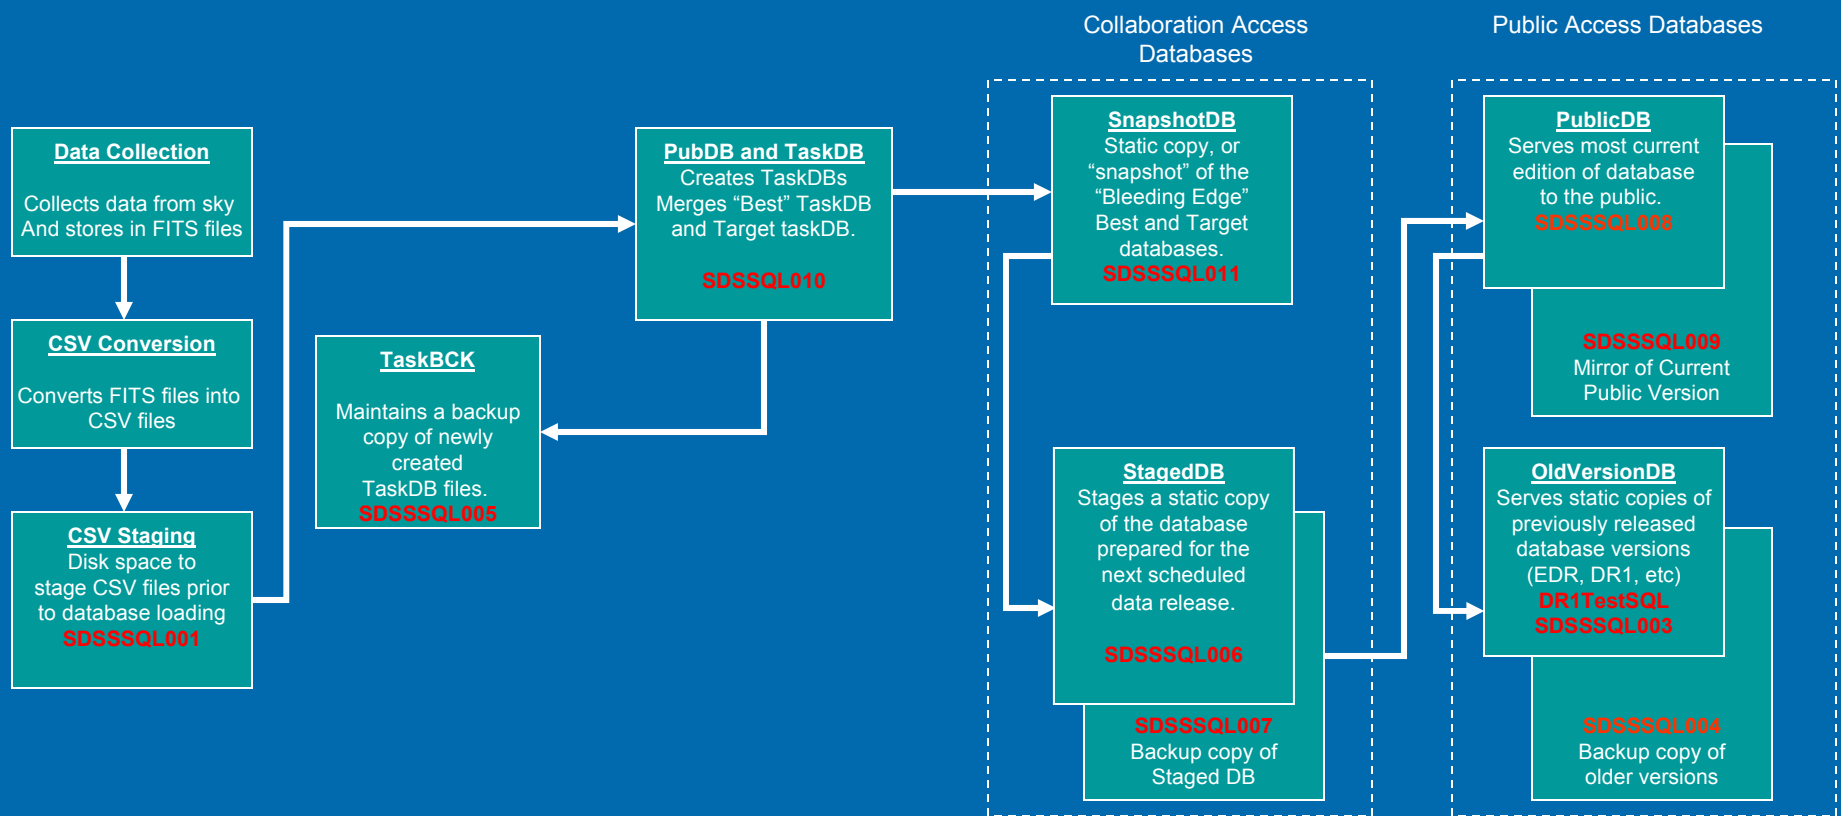
- We propose to buy two new machines viz SDSSSQL008 and SDSSSQL009 in March 2004 to serve as the StagedDB and backup machines for the "finished" copy of DR3.



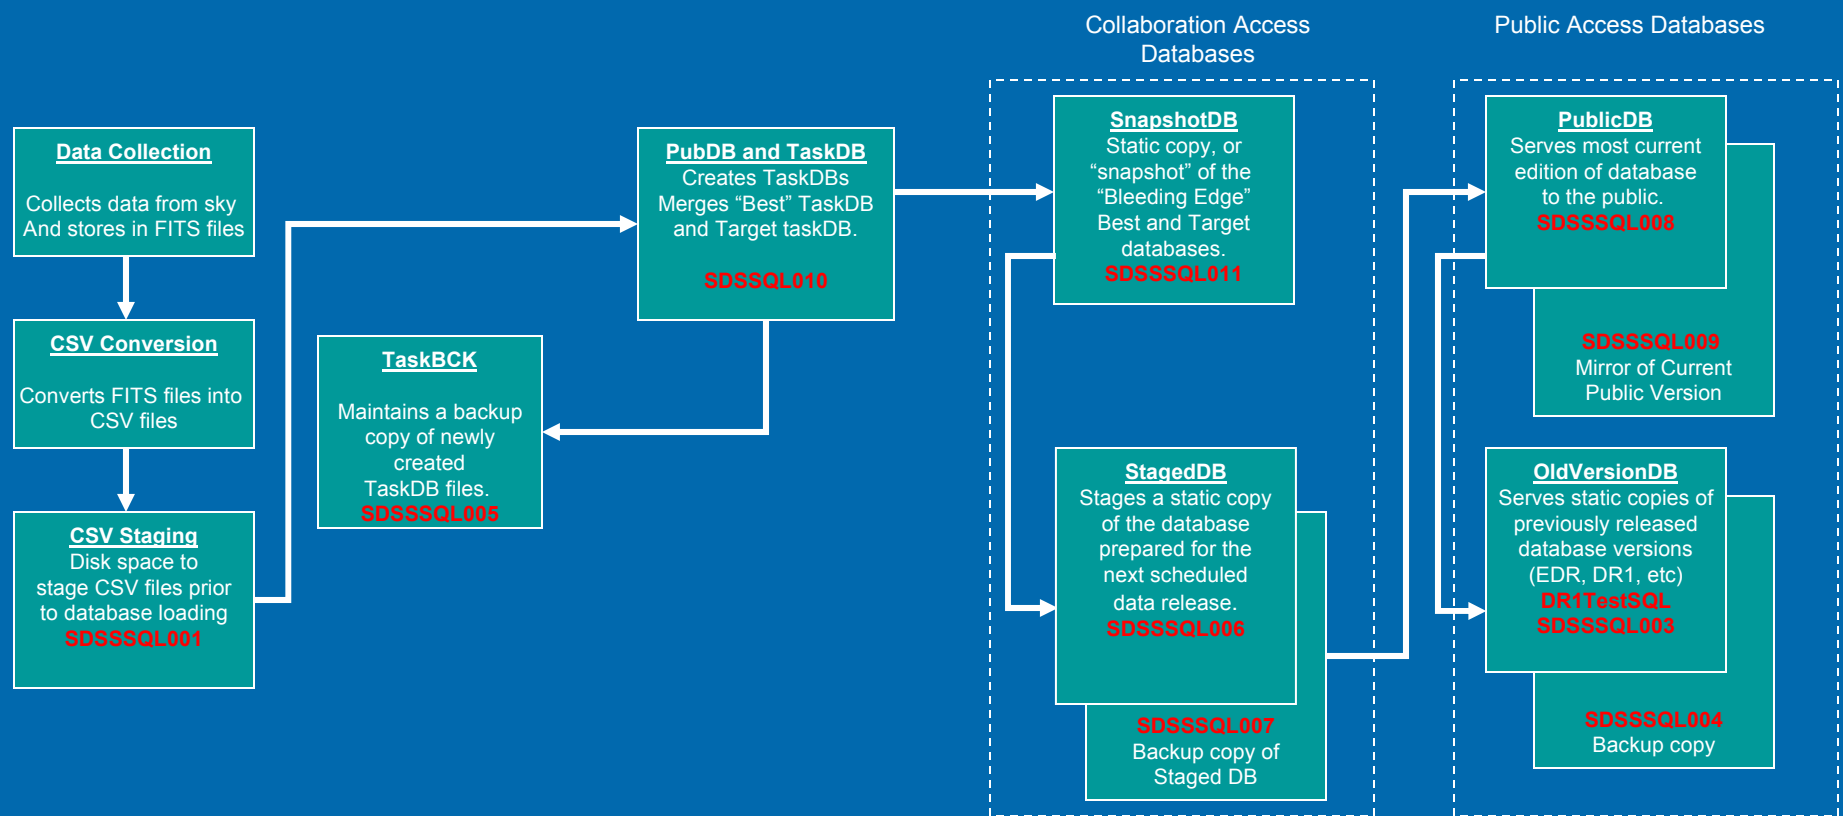
- In Oct 2004, 003 and 004 are pushed down to OldVersionDB
- 008 and 009 now becomes available to public for DR3



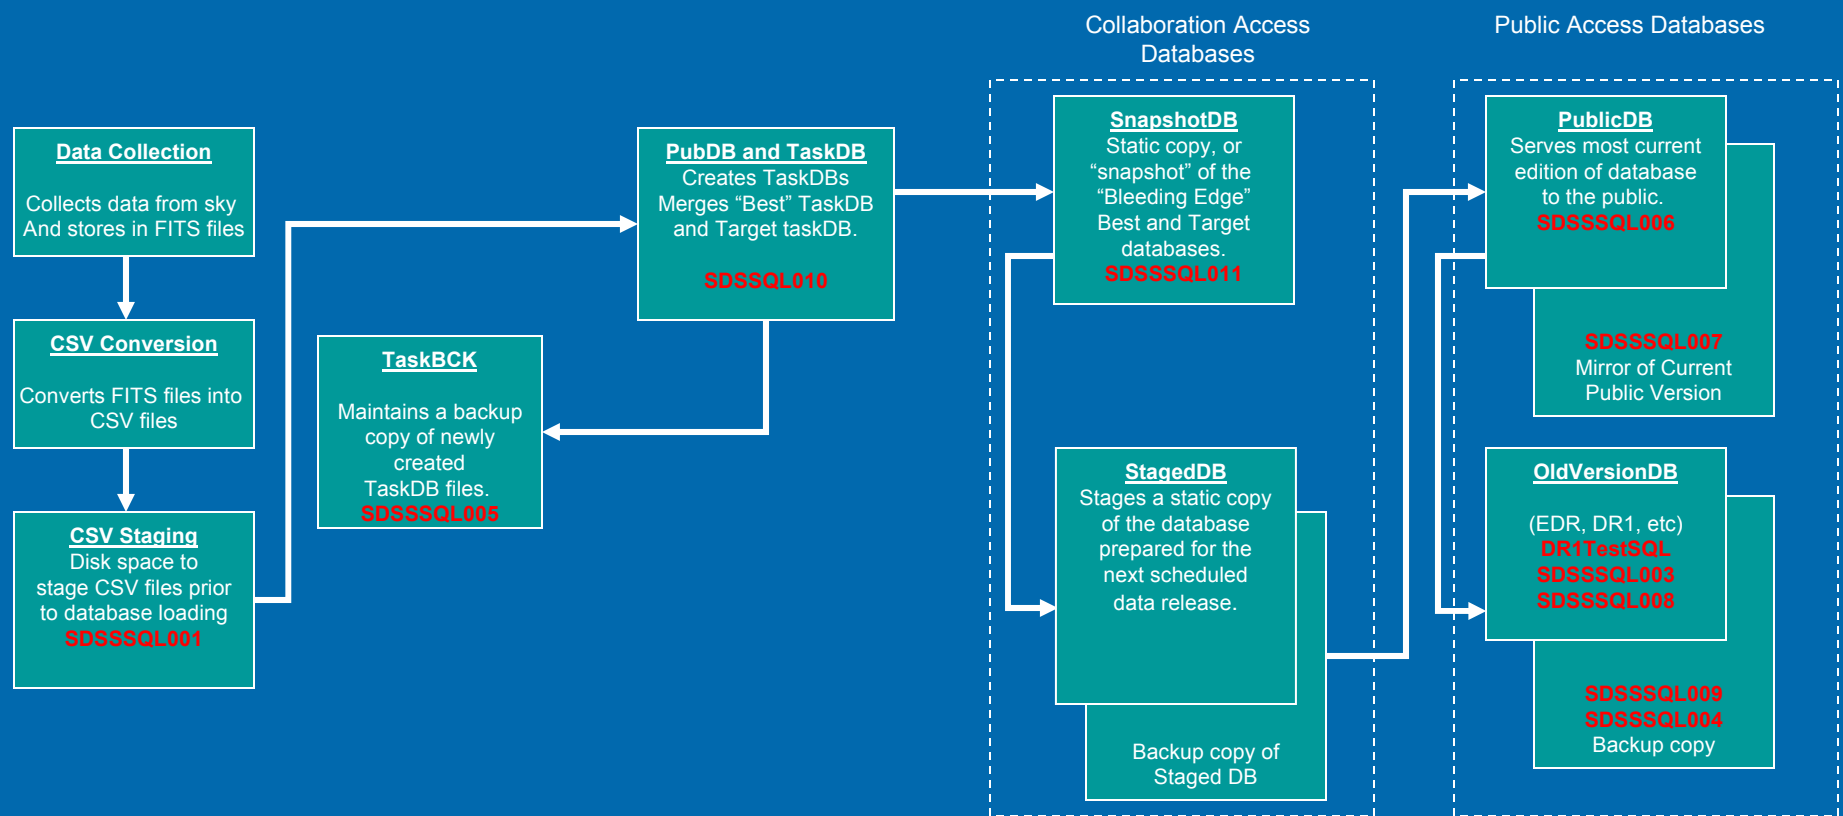
➤ 006 and 007 moves down to become StagedDB and backup StagedDB for DR4



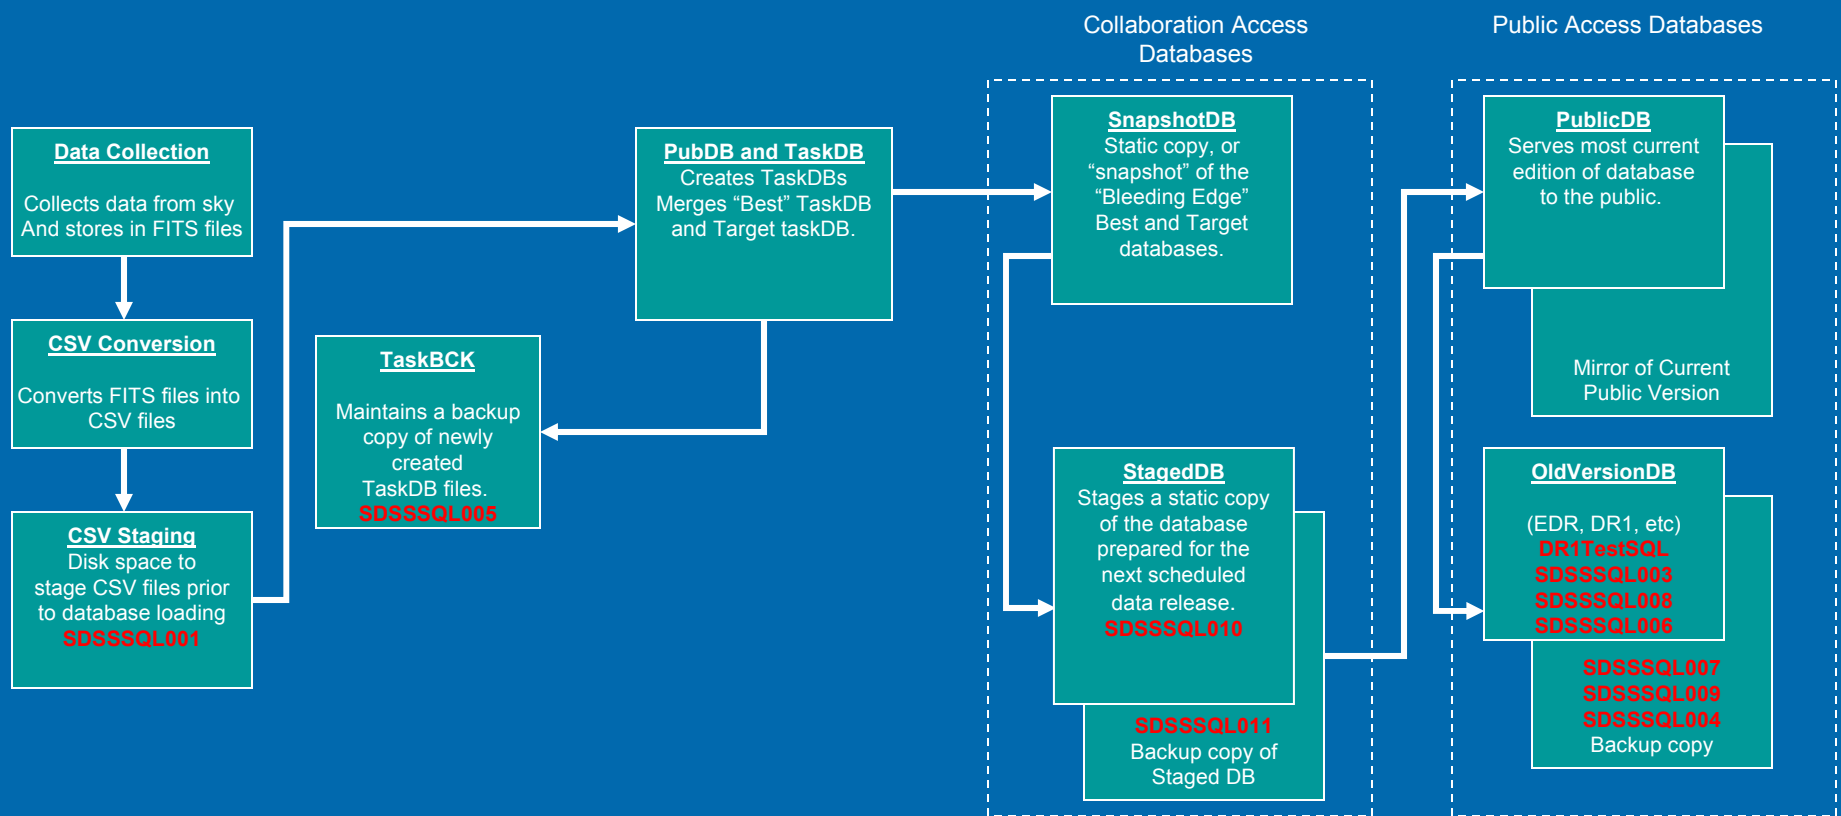
- We propose to buy two new machines viz SDSSSQL010 and SDSSSQL011 in August 2004 to serve as the PubDB and SnapshotDB for DR4.



- In July 2005, 008 and 009 are pushed down to OldVersionDB
- 006 and 007 is made available to public for DR4



- By Sep 2005, 010 and 011 now becomes StagedDB and backup StagedDB for DR5
- 006 and 007 are pushed down to OldVersionDB



- 010 and 011 is made available to public for DR5
- DR5 constitute of Final Data Release for SDSS.
- It is currently scheduled for Oct 2005.